SPECIFICATION

Method, Apparatus, Computer program, Computer system, and
Computer-readable storage medium for Representing and
Searching for an Object in an Image

Technical Field

The present invention relates to the representation of
an object appearing in a still or video image, such as an
image stored in a multimedia database, especially for
searching purposes, and to a method and apparatus for
searching for an object using such a representation.

Background Art

In applications such as image or video libraries, it is
desirable to have an efficient representation and storage of
the outline or shape of objects or parts of objects
appearing in still or video images. A known technique for
shape-based indexing and retrieval uses Curvature Scale
Space (CSS) representation. Details of the CSS
representation can be found in the papers "Robust and
Efficient Shape Indexing through Curvature Scale Space" Proc.
British Machine Vision conference, pp 53-62, Edinburgh, UK,
1996 and "Indexing an Image Database by Shape Content using
Curvature Scale Space" Proc. IEE Colloquium on Intelligent

1

Databases, London 1996, both by F. Mokhtarian, S. Abbasi and J. Kittler, the contents of which are incorporated herein by reference.

The CSS representation uses a curvature function for the outline of the object, starting from an arbitrary point on the outline. The curvature function is studied as the outline shape is evolved by a series of deformations which smooth the shape. More specifically, the zero crossings of the derivative of the curvature function convolved with a family of Gaussian filters are computed. The zero crossings are plotted on a graph, known as the Curvature Scale Space, where the x-axis is the normalised arc-length of the curve and the y-axis is the evolution parameter, specifically, the parameter of the filter applied. The plots on the graph form loops characteristic of the outline. Each convex or concave part of the object outline corresponds to a loop in the CSS image. The co-ordinates of the peaks of the most prominent loops in the CSS image are used as a representation of the outline.

To search for objects in images stored in a database matching the shape of an input object, the CSS representation of an input shape is calculated. The similarity between an input shape and stored shapes is determined by comparing the position and height of the peaks in the respective CSS images using a matching algorithm.

A problem with the known CSS representation is that the peaks for a given outline are based on the curvature function which is computed starting from an arbitrary point on the outline. If the starting point is changed, then there is a cyclic shift along the x-axis of the peaks in the CSS image. Thus, when a similarity measure is computed, all possible shifts need to be investigated, or at least the most likely shift. This results in increased complexity in the searching and matching procedure.

Accordingly the present invention provides a method of representing an object appearing in a still or video image, by processing signals corresponding to the image, the method comprising deriving a plurality of numerical values associated with features appearing on the outline of an object starting from an arbitrary point on the outline and applying a predetermined ordering to said values to arrive at a representation of the outline. Preferably, said values are derived from a CSS representation of said outline, and preferably they correspond to the CSS peak values.

As a result of the invention, the computation involved in matching procedures can be greatly reduced, without a significant reduction in the retrieval accuracy.

3

Disclosure of Invention

A method of representing an object appearing in a still or video image, by processing signals corresponding to the image set forth in claim 1, the method comprises deriving a plurality of numerical values associated with features appearing on the outline of an object starting from an arbitrary point on the outline and applying a predetermined ordering to said values to arrive at a representation of the outline.

In a method set forth in claim 2, the predetermined ordering is such that the resulting representation is independent of the starting point on the outline.

In a method set forth in claim 3, the numerical values reflect points of inflection on the curve.

In a method set forth in claim 4, a curvature scale space representation of the outline is obtained by smoothing the outline in a plurality of stages using a smoothing parameter sigma, resulting in a plurality of outline curves, using values for the maxima and minima of the curvature of each outline curve to derive curves characteristic of the original outline, and selecting the co-ordinates of peaks of said characteristic curves as said numerical values.

In a method set forth in claim 5, the co-ordinates of

4

the characteristic curves correspond to an arc-length parameter of the outline and the smoothing parameter.

In a method set forth in claim 6, the peak co-ordinate values are ordered on the basis of the peak height values, corresponding to the smoothing parameter.

In a method set forth in claim 7, the values are ordered starting from the greatest value.

In a method set forth in claim 8, the values are ordered in decreasing size.

In a method set forth in claim 9, the values are ordered starting from the smallest value.

A method of representing an object appearing in a still or video image, by processing signals corresponding to the image set forth in claim 10, the method comprises deriving a plurality of numerical values associated with features appearing on the outline of an object to represent said outline and deriving a factor indicating the reliability of said representation using a relationship between at least two of said values.

In a method set forth in claim 11, the factor is based on the ratio between two of said values.

In a method set forth in claim 12, the ratio is of the

5

two greatest values.

In a method set forth in claim 13, a curvature scale space representation of the outline is obtained by smoothing the outline in a plurality of stages using a smoothing parameter sigma, resulting in a plurality of outline curves, using values for the maxima and minima of the curvature of each outline curve to derive curves characteristic of the original outline, and selecting the co-ordinates of peaks of said characteristic curves as said numerical values.

In a method set forth in claim 14, the values are derived using a method as claimed in any one of claims 1 to 9.

A method of searching for an object in a still or video image by processing signals corresponding to images set forth in claim 15, the method comprises inputting a query in the form of a two-dimensional outline, deriving a descriptor of said outline using a method as claimed in any one of claims 1 to 9, obtaining a descriptor of objects in stored images derived using a method as claimed in any one of claims 1 to 9 and comparing said query descriptor with each descriptor for a stored object, and selecting and displaying at least one result corresponding to an image containing an

6

object for which the comparison indicates a degree of similarity between the query and said object.

In a method set forth in claim 16, a factor is derived for the query outline and for each stored outline using a method as claimed in any one of claims 10 to 12, and the comparison is made using the predetermined ordering only or the predetermined ordering and some other ordering depending on said factors.

A method of representing a plurality of objects appearing in still or video images, by processing signals corresponding to the images set forth in claim 17, the method comprises deriving a plurality of numerical values associated with features appearing on the outline of each object and applying the same predetermined ordering to said values for each outline to arrive at a representation of each outline.

An apparatus set forth in claim 18 is adapted to implement a method as claimed in any one of claims 1 to 17.

A computer program set forth in claim 19 implements a method as claimed in any one of claims 1 to 17.

A computer system set forth in claim 20 is programmed to operate according to a method as claimed in any one of

claims 1 to 17.

A computer-readable storage medium set forth in claim 21 stores computer-executable process steps for implementing a method as claimed in any one of claims 1 to 17.

A method of representing objects in still or video images set forth in claim 22 is described with reference to the accompanying drawings.

A method of searching for objects in still or video images set forth in claim 23 is described with reference to the accompanying drawings.

A computer system set forth in claim 24 is described with reference to the accompanying drawings.

Brief Description of the Drawings

Fig. 1 is a block diagram of a video database system;

Fig. 2 is a drawing of an outline of an object;

Fig. 3 is a CSS representation of the outline of Fig. 2; and

Fig. 4 is a block diagram illustrating a searching method.

Best Mode for Carrying Out the Invention

First embodiment

Fig. 1 shows a computerised video database system according to an embodiment of the invention. The system includes a control unit 2 in the form of a computer, a display unit 4 in the form of a monitor, a pointing device 6 in the form of a mouse, an image database 8 including stored still and video images and a descriptor database 10 storing descriptors of objects or parts of objects appearing in images stored in the image database 8.

A descriptor for the shape of each object of interest appearing in an image in the image database is derived by the control unit 2 and stored in the descriptor database 10. The control unit 2 derives the descriptors operating under the control of a suitable program implementing a method as described below.

Firstly, for a given object outline, a CSS representation of the outline is derived. This is done using the known method as described in one of the papers mentioned above.

More specifically, the outline is expressed by a representation $\Psi = \{(x(u), y(u), u \in [0, 1]\}$ where u is a normalised arc length parameter.

9

The outline is smoothed by convolving Ψ with an ID Gaussian kernel g(u, σ), and the curvature zero crossings of the evolving curve are examined as σ changes. The zero crossing are identified using the following expression for the curvature:

$$k(u,\sigma) = \frac{X_u(u,\sigma)Y_{uu}(u,\sigma) - X_{uu}(u,\sigma)Y_u(u,\sigma)}{(X_u(u,\sigma)^2 + Y_u(u,\sigma)^2)^{3/2}}$$

where

$$X(u,\sigma) = x(u) * g(u,\sigma) \qquad Y(u,\sigma) = y(u) * g(u,\sigma)$$

and

$$X_u(u,\sigma) = x(u) * g_u(u,\sigma) \qquad X_{uu}(u,\sigma) = x(u) * g_{uu}(u,\sigma)$$

In the above, * represents convolution and subscripts represent derivatives.

The number of curvature zero crossings changes as σ changes, and when σ is sufficiently high Ψ is a convex curve with no zero crossings.

The zero crossing points (u, σ) are plotted on a graph , known as the CSS image space. This results in a plurality of curves characteristic of the original outline. The peaks of the characteristic curves are identified and the corresponding co-ordinates are extracted and stored. In general terms, this gives a set of n co-ordinate pairs

[(x1,y1), (x2,y2), ...(xn,yn)], where n is the number of peaks, and xi is the arc-length position of the ith peak and yi is the peak height.

The order and position of characteristic curves and the corresponding peaks as they appear in the CSS image space depends on the starting point for the curvature function described above. According to the invention, the peak co-ordinates are re-ordered using a specific ordering function.

Ordering is performed by a one-to-one mapping T of the peak indices {1 ...n}to a new set of indices {1 ...n}.

In this embodiment, the co-ordinate pairs are ordered by considering the size of the y co-ordinates. Firstly, the highest peak is selected. Suppose the kth peak is the most prominent. Then (xk, yk) becomes the first in the ordered set of values. In other words, $T(k) = 1$. Similarly, the other peak co-ordinates are re-ordered in terms of decreasing peak height. If two peaks have the same height, then the peak having the x-co-ordinate closest to that of the preceding co-ordinate pair is placed first. In other words, each co-ordinate pair having an original index i is assigned a new index j where $T(i) = j$ and $y_j >= y(j+1)$. Also, each value xi is subjected to a cyclic shift of -xk.

11

As a specific example, the outline shown in Fig. 2 results in a CSS image as shown in Fig. 3. Details of the co-ordinates of the peaks of the curves in the CSS image are given in Table 1 below.

Peak Index	X	Y
1	0.124	123
2	0.68	548
3	0.22	2120
4	0.773	1001
5	0.901	678

Table 1.

The peaks are ordered using the ordering described above. In other words, the co-ordinates are ordered in terms of decreasing peak height. Also, the x co-ordinates are all shifted towards zero by an amount equal to the original x co-ordinate of the highest peak. This results in re-ordered peak co-ordinates as given in Table 2 below.

Peak Index	X	Y
1	0	2120
2	0.553	1001
3	0.681	678
4	0.46	548
5	0.904	123

Table 2.

These re-ordered peak co-ordinates form the basis of

the descriptor stored in the database 10 for the object outline. In this embodiment, the peak co-ordinates are stored in the order shown in Table 2. Alternatively, the co-ordinates can be stored in the original order, together with an associated indexing indicating the new ordering.

Second embodiment

An alternative method of representing the object outline according to a second embodiment will now be described.

A CSS representation of the outline is derived as described above. However, the ordering of the peak co-ordinates is different from the ordering in Embodiment 1 described above. More specifically, firstly the highest peak is selected. Suppose peak k is the most prominent one. Then (x_k, y_k) becomes the first peak in the ordered set of peaks. The subsequent peaks are ordered so that for peak co-ordinates of original index i, then $T(i) = j$, and $x_j <= x(j+1)$. Also, all values x_i are shifted downwards by an amount x_k equal to the original x co-ordinate of original peak k.

In other words, in the ordering method according to embodiment 2, the highest peak is selected and placed first,

13

and then the remaining peaks follow in the original sequence starting from the highest peak.

Table 3 below shows the peak values of Table 1 ordered according to the second embodiment.

Peak Index	X	Y
1	0	2120
2	0.46	548
3	0.553	1001
4	0.681	678
5	0.904	123

Table 3.

In a development of embodiments 1 and 2 described above, a confidence factor (CF) is additionally associated with each representation of a shape. The CF is calculated from the ratio of the second highest and the highest peak values for a given shape.

For the outline shown in Fig. 2, the CF value is CF = 1001/2120. In this example, the CF is quantized by rounding to the nearest 0.1 to reduce storage requirements. Accordingly, here CF = 0.5.

The CF value in this example is a reflection of the accuracy or uniqueness of the representation. Here, a CF value close to one means low confidence and a CF value close to zero means high confidence. In other words, the closer are the two highest peak values, the less likely it is that the representation is accurate.

The CF value can be useful when performing a matching

15

procedure, as will be shown in the following description.

Third embodiment

A method of searching for an object in an image in accordance with an embodiment of the invention will now be described with reference to Fig. 4 which is a block diagram of the searching method.

Here, the descriptor database 10 of the system of Fig. 1 stores descriptors derived according to the first ordering method described above together with associated CF values.

The user initiates a search by drawing an object outline on the display using the pointing device (step 410). The control unit 2 then derives a CSS representation of the input outline and orders the peak co-ordinates in accordance with the same ordering function used for the images in the database to arrive at a descriptor for the input outline (step 420). The control unit 2 then also calculates a CF value for the input outline by calculating the ratio of the second highest peak value to the highest peak value and quantizing the result (step 430).

The control unit 2 then compares the CF value for the input outline with a predetermined threshold (step 440). In this example, the threshold is 0.75. If the CF value is

lower than the threshold, indicating a relatively high confidence in the accuracy of the input descriptor, then the next step is to consider the CF value for the model (ie image stored in the database) under consideration. If the model CF is also lower than the threshold (step 450), then the input and model are compared using the respective descriptors in the predetermined ordering only (step 460). If CF for either the input or the model is greater than the threshold, then matching is performed by comparing all possible different orderings of the co-ordinate values in the input descriptors with the model descriptor in the database (step 470).

The matching comparison is carried out using a suitable algorithm resulting in a similarity measure for each descriptor in the database. A known matching algorithm such as described in the above-mentioned papers can be used. That matching procedure is briefly described below.

Given two closed contour shapes, the image curve Ψi and the model curve Ψm and their respective sets of peaks $\{(xi1,yi1),(xi2,yi2),...,(xin,yin)\}$ and $\{(xm1,ym1), (xm2,ym2), ...,(xmn,ymn)\}$ the similarity measure is calculated. The similarity measure is defined as a total cost of matching

of peaks in the model into peaks in the image. The matching
which minimises the total cost is determined using a dynamic
programming. The algorithm recursively matches the peaks
from the model to the peaks from the image and calculates
the cost of each such match. Each model peak can be matched
with only one image peak and each image peak can be matched
with only one model peak. Some of the model and or image
peak may remain unmatched, and there is an additional
penalty cost for each unmatched peak. Two peaks can be
matched if their horizontal distance is less then 0.2. The
cost of a match is the length of the straight line between
the two matched peaks. The cost of an unmatched peak is its
height.

In more detail the algorithm works by creating and
expanding a tree-like structure, where nodes correspond to
matched peaks:

1.   Create starting node consisting of the largest
maximum of the image (xik, yik) and the largest maximum of
the model (xir,yir).

2.   For each remaining model peak which is within 80
percent of the largest maximum of the image peaks create an
additional starting node.

3. Initialise the cost of each starting node created in 1 and 2 to the absolute difference of the y-coordinate of the image and model peaks linked by this node.

4. For each starting node in 3, compute the CSS shift parameter alpha, defined as the difference in the x (horizontal) coordinates of the model and image peaks matched in this starting node. The shift parameter will be different for each node.

5. For each starting node, create a list of model peaks and a list of image peaks. The list hold information which peaks are yet to be matched. For each starting node mark peaks matched in this node as "matched", and all other peaks as "unmatched".

6. Recursively expand a lowest cost node (starting from each node created in steps 1-6 and following with its children nodes) until the condition in point 8 is fulfilled. To expand a node use the following procedure:

7. Expanding a node:

If there is at least one image and one model peak left unmatched:

select the largest scale image curve CSS maximum which is not matched (xip,yip). Apply the starting node

shift parameter (computed in step 4) to map the selected

maximum to the model CSS image - now the selected peak has

coordinates (xip-alpha, yip). Locate the nearest model curve

peak which is unmatched (xms,yms). If the horizontal

distance between the two peaks is less then 0.2 (i.e: |xip-

alpha- xms | < 0.2), match the two peaks and define the cost

of the match as the length of the straight line between the

two peaks. Add the cost of the match to the total cost of

that node. Remove the matched peaks from the respective

lists by marking them as "matched". If the horizontal

distance between the two peaks is greater than 0.2, the

image peak (xip,yip) cannot be matched. In that case add its

height yip to the total cost and remove only the peak

(xip,yip) from the image peak list by marking it as

"matched".

Otherwise (There are only image peaks or there are

only model peaks left unmatched):

Define the cost of the match as the height of the

highest unmatched image or model peak and remove that peak

from the list.

8.   If after expanding a node in 7 there are no

unmatched peaks in both the image and model lists, the

matching procedure is terminated. The cost of this node is the similarity measure between the image and model curve. Otherwise, go to point 7 and expand the lowest cost node.

The above procedure is repeated with the image curve peaks and the model curve peaks swapped. The final matching value is the lower of the two.

As another example, for each position in the ordering, the distance between the input x value and the corresponding model x value and the distance between the input y value and the corresponding model y value are calculated. The total distance over all the positions is calculated and the smaller the total distance, the closer the match. If the number of peaks for the input and the model are different, the peak height for the leftovers is included in the total distance.

The above steps are repeated for each model in the database (step 480).

The similarity measures resulting from the matching comparisons are ordered (step 490) and the objects corresponding to the descriptors having similarity measures indicating the closest match (i.e. here the lowest similarity measures) are then displayed on the display unit

4 for the user (step 500). The number of objects to be displayed can be pre-set or selected by the user.

In the above embodiment, if the CF value is greater than the threshold, then all possible orderings of the input descriptor values are considered in the matching. It is not necessary to consider all possible orderings, and instead only some possible orderings may be considered, such as some or all cyclic shifts of the original CSS representation. Furthermore, in the above embodiment, the threshold value is set to 0.75, but the threshold can be set to different levels. For example, if the threshold is set to zero, then all matches are performed by analysis of some or all possible orderings. This increases the amount of computation required compared with case when threshold is above zero, but since the peaks have already been ordered and their x-coordinate adjusted for a particular starting point or object rotation, the amount of computation required is reduced compared with the original system where no such adjustment has been made. Consequently, by setting the threshold to zero the system offers some reduction in computational cost and the retrieval performance is exactly the same as in the original system.

Alternatively, if the threshold is set to one, then matching is performed using only the stored ordering. There is then a significant reduction in computation required, with only a small deterioration in retrieval accuracy.

Various modifications of the embodiments described above are possible. For example, instead of ordering the CSS peak co-ordinate values as described in embodiments 1 and 2 other orderings can be used. For example, the values can be placed in order of increasing rather than decreasing peak height. Instead of storing the ordered values in the database, the ordering can be carried out during the matching procedure.

Industrial Applicability

A system according to the invention may, for example, be provided in an image library. Alternatively, the databases may be sited remote from the control unit of the system, connected to the control unit by a temporary link such as a telephone line or by a network such as the internet. The image and descriptor databases may be provided, for example, in permanent storage or on portable data storage media such as CD-ROMs or DVDs.

Components of the system as described may be provided in software or hardware form. Although the invention has been described in the form of a computer system, it could be implemented in other forms, for example using a dedicated chip.

Specific examples have been given of methods of representing a 2D shape of an object and of methods for calculating values representing similarities between two shapes but any suitable such methods can be used.

The invention can also be used, for example, for matching images of objects for verification purposes, or for filtering.